Table-top Computed Lighting for Practical Digital Photography

Ankit Mohan, Reynold Bailey, Jonathan Waite, Jack Tumblin, *Member, IEEE*, Cindy Grimm, *Member, IEEE*, Bobby Bodenheimer, *Senior Member, IEEE*

***Abstract*— We apply simplified image-based lighting methods to reduce the equipment, cost, time, and specialized skills required for high-quality photographic lighting of desktop-sized static objects such as museum artifacts. We place the object and a computer-steered moving-head spotlight inside a simple foam-core enclosure, and use a camera to record photos as the light scans the box interior. Optimization, guided by interactive user sketching, selects a small set of these photos whose weighted sum best matches the user-defined target image. Unlike previous image-based re-lighting efforts, our method requires only one light source, yet can achieve high resolution light positioning to avoid multiple sharp shadows. A reduced version uses only a hand-held light, and may be suitable for battery-powered, field photography equipment that fits into a backpack.**

***Index Terms*— Enclosure Lighting, Handheld Lighting, Controllable Lighting, Digital Photography**

I. Introduction

MODERN digital cameras have made picture-taking easier and more interactive. However lighting a scene for good photography is still difficult, and practical methods to achieve good lighting have scarcely changed. We show that sketch-guided optimization and simplified forms of image-based lighting can substantially reduce the cost, equipment, skill, and patience required for small-scale studio-quality lighting.

Good studio lighting is difficult because it is a 4D inverse problem that photographers must solve by making successive approximations guided by years of experience. For non-experts, good studio lighting can be surprisingly frustrating. Most people can specify the lighting they want in screen space (e.g. "get rid of this obscuring highlight, make some shadows to reveal rough texture here, but fill in the shadows there"), but determining what kind of lights to use, where to place them, and how to orient them is difficult.

We are interested in camera-assisted lighting for human-scale, desktop-sized static objects. We want lighting that accurately reveals the shape, texture, materials, and most visually meaningful features of the photographed item. In particular, we envision two applications of our work. The first is a method to help museum curators as they gather digital photographic archives of their vast collections of items. The second application allows users with ordinary photographic equipment to make appealing photographs of items that might be displayed, for example, on such sites as Ebay. We believe this latter application is especially appropriate since users will be motivated to make appealing, but physically-based, photographs of items for sale. A pictorial synopsis of our system is shown in Figure 1.

Depending on the application, then, data for our process can be acquired in two ways, automatically and manually. The automatic method, the one we envision for use by museum curators, is more flexible and convenient but requires additional hardware over the manual method, which we envision as the more commodity-level method. Both techniques have merit, yet the challenges of the automatic method are more interesting, and so our discussion is slanted toward the automatic method. However, both methods are described in detail.

Pioneering work in image-based lighting ([8], [9], [12], [17]) offers promising approaches that can help with the photographic lighting problem. Unfortunately, most of these require too many precise measurements and adjustments for day-to-day use outside the laboratory. Precision is required to address more ambitious goals such as recovering shape, BRDF, and appearance under arbitrary viewing and lighting conditions. For the much smaller, yet more widespread problem of photographic lighting, we need far less: we need a method that requires less time, expense, and complexity, yet allows users who are not lighting experts to quickly find the lighting they want.

This paper offers three contributions. We extend existing image-based lighting ideas to reduce the required equipment to a single light source and single camera; we replace trial-and-error light repositioning with optimization and on-screen painting; and we reduce the need for large amounts of high dynamic range photography, thus reducing the capture time. The result is a novel and inexpensive system that a novice can use to intuitively describe and obtain the desired lighting for a photograph.

II. Related Work

Lighting has long been recognized as a hard problem in computer graphics and many papers have explored optimization for light placement and other parameters ([7], [14],

Manuscript received YY/YY/2006

A.Mohan and J. Tumblin are with Department Of Electrical Engineering and Computer Science, Northwestern University, 3.320 Ford Design Center, 2133 Sheridan Road, Evanston, IL 60208, USA. E-mail: {ankit,jet}@cs.northwestern.edu

R. Bailey and C. Grimm are with Department of Computer Science and Engineering, Campus Box 1045, Washington University in St. Louis, One Brookings Drive, St. Louis, MO 63130, USA. E-mail: {rjb1,cmg}@cse.wustl.edu

J. Waite and B. Bodenheimer are with Department of Electrical Engineering and Computer Science, Vanderbilt University, VU Station B #351679, 2301 Vanderbilt Place, Nashville, TN 37235, USA. E-mail: {jon.waite,bobby.bodenheimer}@vanderbilt.edu.

(a) Good result using professional lighting; (b) which is complicated and slow. (c) Our lighting results; (d) with a steerable spotlight and an uncalibrated enclosure.

Fig. 1. Our system gives results comparable to professional lighting even when used by novice users.

[21], [24], [25]). Some of these systems also used painting interfaces to specify desired lighting in screen space ([21], [22], [24]), and we use a similar approach to make lighting for photography more intuitive. The system by Shacked et al. [25] was even able to provide fully automatic lighting by applying image quality metrics. However, all of these systems require 3D information unavailable in our photographic application.

As shown in Figures 1(d) and 3, our system in its most automatic form uses an enclosure. Several commercially available photographic enclosures exist, but they are primarily used to achieve very soft lighting; they do not help users solve light placement problems. These systems include diffusive tents [20], photo-boxes [18], and translucent back-lit platforms with an array of individually dimmed light sources [4].

Image-based methods have also been used to perform arbitrary relighting of well-measured objects. Most methods, including ours, perform relighting using a weighted sum of differently lit basis images, done first by Nimeroff et al. [19]. The key observation of this work is that light and materials interact linearly. However, prior efforts used more elaborate and expensive equipment because their goals were different from ours. These included measurement of a 4D slice of the reflectance field of the human face by Debevec et al. [8], museum artifacts measured by a rotating-arm light stage by Hawkins et al. [12], an ingenious but expensive system by Debevec et al. [9] for real-time video playback and measurement of light fields, a dome of electronic flashes for real time image relighting by Malzbender et al. [15], a free form light stage to enable portable gathering of light-field data with some calibration by Masselus et al. [16], and full 4D incident light measurements by Masselus et al. [17]. In all of these cases, data-gathering required either customized equipment or collection times much longer than would be practical for photographic lighting.

Three recent systems also offered novel sketch guided relighting from basis images. Akers et al. [2] used a robotic light-positioning gantry to gather precisely lit images, and like us, also provided a painting interface to guide re-lighting, but unlike us they used spatially varying weights that could produce physically impossible lighting. Agarwala et al. [1] used sketch-guided graph-cut segmentation coupled with gradient domain fusion to seamlessly merge several photographs. They demonstrated merging differently lit photographs to create novel illumination conditions. Though their interaction scheme worked well for a small number of images($\sim$10), it may be impractical for the hundreds of images required for complete control over lighting directions. Also, their system does nothing to help the user with light placement, and may produce physically unrealizable results. Anrys and Dutre [3] used a Debevec-style light stage with $\sim$40 fixed, low powered light sources and a painting interface to guide lighting. Their optimization only found light intensities, and light placement is still left up to the user. Also, their point light sources can cause multiple shadows and highlights which may be undesirable for archival purposes. The data capture time is also high since they capture high dynamic range photos for every light location.

Our work is different from these previous systems in that our system does not force users to decide on correct or complete light source placement. This result is possible because our capture process is significantly different from prior methods, and better suited for the task of photography. We require less than five minutes to complete the initial image capture and a few more minutes to get the final result. The equipment required is minimal and portable, and our hand-held version can be carried in a backpack. Most similar to our goals and methods is the work of Fuchs et al. [11], who manually swept a light source over walls with an object placed near a probe object. They then used Bayesian techniques to relight the scene. While their method is well-suited to relighting the scene, the end result typically contains more noise than ours.

III. Simplifications and Constraints: 2D lighting

Our goal is to create a system that provides computational help for doing what a good photographer does. In particular, our goal is to provide a photograph of a scene lit in a specific way. Photographers make consistent choices about which types of lights to use, how to adjust them, and where to place them. We show how our streamlined image-based method follows these same choices (see Figure 1); we are not seeking to build a calibrated 4D data set to reconstruct all forms of illumination.

The observation from [19] that lights and materials interact linearly means that if a fixed camera makes an image I_i from a fixed scene lit only by a light L_i, then the same scene lit by many lights scaled by weights w_iL_i will make an image $I_{out} = \sum_i w_i I_i$. Adjusting weights lets us "re-light" the image, as if the weights modulate the lights rather than the

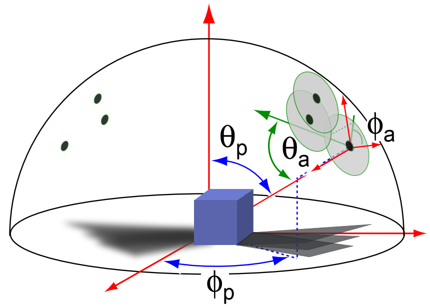

Fig. 2. All possible lighting angles parameterized by light position (θ_p, ϕ_p) and direction (θ_a, ϕ_a). Point light sources (on the left) result in multiple hard shadows, while overlapping area light sources can be used to simulate a larger light source.

images. As we collect more images I_i, we can simulate more lighting possibilities.

How many images do we need to gather? We only need enough images to span the kind of lighting a skilled photographer might explore to get good results in a photo studio. Several common practices in studio lighting guide us.

First, professional photographers choose lamps with broad, nearly uniform beams of light, often with a reflector and lens to help direct more light forward. Second, they adjust light placement angles carefully, but not their distances. Distance to the light affects foreshortening of shadow shapes, but these effects are subtle and rarely noticed in still images. Third, they adjust lights to control shadow softness versus sharpness. Light sources (or more accurately, the shadows they form) become 'softer' by increasing the angular extent as measured from the lit object. Fourth, they seek out lighting arrangements that produce a simple set of shadows and highlights that best reveal the object's shape, position, and surface qualities. They avoid complex overlapped shadows, lack of shadows due to overly-soft light, and contrast extremes from large specular highlights or very dark shadows. Simpler shadows usually mean fewer lights, and thus fewer basis images.

Accordingly, we use commercially available light sources instead of custom or special-purpose devices. We place light sources at a moderate distance (typically ~ 1 meter). We use small-to-moderate area 'soft' light sources instead of the much sharper point-like sources often used in earlier approaches. Overlapped soft shadows blend far less noticeably than sharp shadows from the same light positions (as shown in Figure 2), so that soft lights require us to gather fewer images to avoid multiple shadow artifacts. Also, overlapping area light sources can be combined to produce a larger area light source.

Note that we do not need to know the light positions or their absolute intensities for our images; we select weights w_i and images I_i by their ability to match the lighting target images a user sketches for us. Although we do not need calibration, we can make use of the ability to return the light to a previous position. This feat is possible using a single, commercially available steerable light. We also require consistency in the light response curve, available in commercial digital cameras.

We avoid the use of high-dynamic-range (HDR) photographs where possible, as these typically require multiple calibrated exposures and computation to merge them [10]. Instead, we rely on the camera's automatic exposure adjustments to capture images suitable for interactive lighting design. We only resort to HDR capture methods for basis images with large over-exposed regions. Under-exposed regions can be ignored, as their contributions are already invisible.

Formally, arbitrary external illumination is 4D for a desktop scene: $L(\theta_p, \phi_p, \theta_a, \phi_a) = L(\Theta)$. Suppose that the photographed object receives all of its light from a hemisphere of tiny, invisible, inward-pointing video projectors at radius r. Each projector's position in desktop polar coordinates is (θ_p, ϕ_p). Each projector's center-most pixel $P(\theta_a = 0, \phi_a = 0)$ forms a ray that illuminates the center point of our desktop, and in the projector's polar coordinates the other pixels are $P(\theta_a, \phi_a)$, as shown in Figure 2. All of the combined projectors' light output is the 4D incident light field, and describes all possible lighting. To simulate all possible lighting, we would need a new image I_i to capture light from each pixel of each video projector. Instead, we use only broad beams of light (e.g. $P(\theta_a, \phi_a) \cong \cos(\theta_a) cos(\phi_a)$), regular sampling of light placement angles (θ_p, ϕ_p), and specify 'softer' to 'sharper' shadows by varying the angular extent (θ_p, ϕ_p) as measured from the lit object. This angular extent should not be confused with the lamp's beam width (θ_a, ϕ_a); in our 'hemisphere of video projectors' analogy, beam width sets the image from a projector, but angular extent sets the number of adjacent projectors that emit this same image.

In summary, rather than recreate arbitrary 4D incident light fields, we use weighted sums of basis images that represent the type of lighting used by professional photographers. This method is much more practical and efficient, with little, if any, loss of generality.

IV. Method - Light-Gathering

We construct a high quality user-guided picture in three steps. First, the system or user captures a set of photos for densely sampled lighting angles for the photographed object. Second, the user iteratively paints the desired lighting by simple lighten-darken operations to generate a target image. The system finds weights w_i for each photo such that their weighted sum matches the target image in the least squares sense. A weighted sum of these images gives the final result. If the result is not satisfactory, the user can sketch on the current result to create the next iteration's target image.

A. Enclosed Light Source and Image Acquisition

Freed from photometric and angular calibration requirements as discussed in Section III, we are able to build a simple and cost-effective controlled light source. We place the object and a gimbal-mounted moving-head spotlight inside an enclosure of almost any convenient size, shape and material. The powerful computer-aimed light pivots to any desired pan and tilt angle with good repeatability ($\leq \pm 0.5°$) to light any desired spot inside our enclosure. The enclosure acts as a reflector, and effectively provides a controllable 2D area light source around the object. The size and shape of the enclosure is almost irrelevant as long as the light is close enough to the

Fig. 3. The disco-light setup. The object and disco light are both enclosed in a white foam box, with the camera looking in through a window in the enclosure wall farthest from the light.

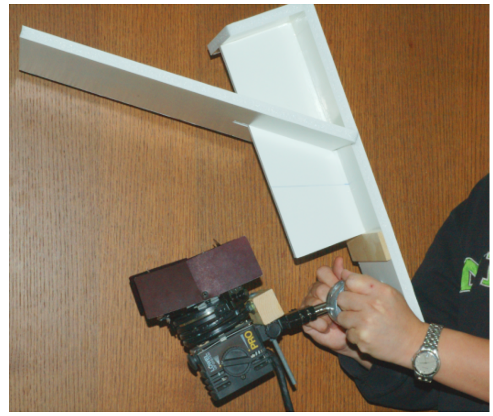

Fig. 4. Hand-held data gathering using a light with an attached foam-core diffuse reflector.

object to keep parallax low, and the light is powerful enough for the camera to get a reasonable exposure.

We built a $1 \times 1 \times 1.5\text{m}^3$ sized box of white $1/2$” foam-core board as our enclosure, and chose an inexpensive moving-head spotlight. The 150-watt *American DJ Auto Spot 150* disco-light, shown in Figure 3 can tilt $270°$, pan $540°$, and includes 9 color filters, gobos and several other fun features. Computer control by the DMX512 protocol is easy to program with the SoundLight USB DMX controller and API. Our foam-core enclosure resembles a hemi-cube on a pair of tables. We place the gimbal light on a small table that lowers its rotation center to the plane of an adjacent taller table holding the photographed object, as shown in Figure 3. Adjacent, but separate, tables reduces vibration, permits gimbal angles to approximate hemisphere angles, and separates the object from the swiveling lamp. We place the camera behind a small opening cut in the enclosure wall on the end farthest from the light source.

The system gathers images rapidly and automatically. Through the DMX512 controller we direct the gimbal light to scan the upper hemisphere of light aiming directions in equal-angle increments as we record individual computer-triggered photographs using auto-exposure. We are able to record hundreds of individual images per minute, and can complete all of the data gathering in about twenty minutes using a Pentium 2GHz computer and a Canon Powershot G3 camera.

To the best of our knowledge, no other image-based lighting work exploits these movable or controllable lights. Enclosed pivoting lights retain many advantages of the more sophisticated lighting systems, avoid multiple sharp shadows, can offer variable “softness” by spot size adjustment, and are much simpler and cheaper to construct. Of course, they do not easily provide accurate lighting direction calibration or point-light illumination, but these features are not needed for our goals.

After recording, we linearize each captured frame (RGB) by applying the camera’s inverse response curve, recovered by the method of Debevec et al. [10], and converted to luminance values. Linear response ensures weighted sums of whole images are accurate representations of physically realizable lighting. We then down-sample the image dataset to 320×240 or smaller for use in the optimization step that follows.

B. Portable, Hand-held Method

Even a foam-core box and a moving-head spotlight are impractical to carry around everywhere. However, the “Free-form light-stage” [16] showed that it is possible to gather calibrated image sets suitable for 2D re-lighting with nothing more than four small light-probe-like spheres, a digital camera and tripod, a hand-held point-light source, possibly battery-powered, and approximately 30 minutes of time to take several hundred digital photographs. While it meets the ambitious goal of incident light field capture, the method would tax anyone’s patience to record more than just a few items. We present a faster and simpler variant that serves our purpose better.

In the method previously described, we require repeatable light source positioning in order to reacquire any needed HDR images. However, if we either ignore over-exposed specular highlights or record high dynamic range images when needed, then *repeatability is not needed*; we can use a hand-held light source in our method as well. As shown in Figure 4, we use a small 250W hand-held light intended for television news cameras, attached to a diffuse reflector (foam core again), and limit the beam width with barn-doors to form a well-defined area light source.

To gather all photos, we hold the light outstretched and move the light on a hemisphere centered about the object in an “image capture dance.” We sample the hemisphere of lighting directions by a polar-coordinate scan in ϕ-first order as the camera takes sequential photographs. A Nikon D70 camera takes a steady stream of photos at about 3 frames/second using autoexposure for each frame. The user stands facing the object, and holds the light at arms’ length while moving the lamp in an arc that passes directly over the object. The user moves the lamp from one side of the table to the other, scanning by π radians in θ axis with constant ϕ, and the natural alignment of their shoulders helps aim the light’s centerline directly at the object. After each pass over the object with the light, the user steps sideways to change the ϕ angle for the next scan, and makes enough of these passes to cover $0 \leq \phi < \pi$ radians. In

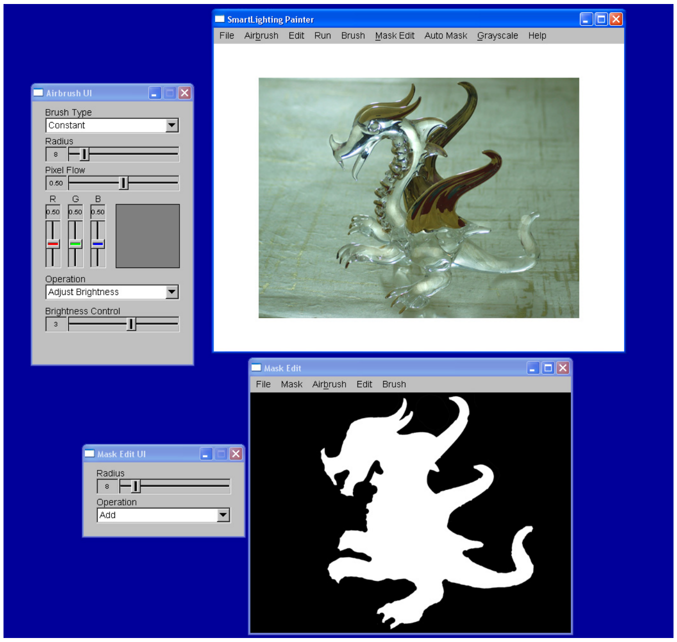

Fig. 5. A screen shot of the user interface for creating target lighting conditions. The target image that the user sketches on is the top image, shown in color. The panel would show the result of any previous iterations as well. The bottom black and white image is the mask; it is typically not displayed unless the user so desires. The control boxes on the left govern what type of adjustment to the target image and mask are being made, e.g., contrast or dodge-and-burn. The control panel also governs the brush size, and other controls for optimization, output assembly, and saving.

practice the user can be more careless with the light, as long as the hemisphere of light is well-sampled and the images are not over-exposed. After the image capture dance is complete, we downsample all images and proceed with the sketch-guided lighting design as before.

We find this process is quite simple and pleasing, and in less than four minutes we can gather a high-quality basis photo set of 120-150 images. An experienced user might not need to scan the whole hemisphere, but can quickly illuminate just the regions where they know they need computed light sources.

V. Method - User Interface and Optimization

A. *Sketch-Guided Lighting Optimization*

After gathering aiming images, users can iteratively specify and refine the lighting by sketching on a *target intensity image*. This gray-scale image (examples in Figure 8) approximates the final output image the user would like to see. The user interface for this application is shown in Figure 5.

For editing the target image, the user starts off either with a simple luminance image of the target, a gray wash (e.g., uniform gray, light gray fading to dark gray across the image), or the previous iteration's result. The user then carries out a series of lighten or darken operations in the different regions of the image to approximate the desired results. The user can paint areas of luminance on the image directly, use a dodge/burn control, or adjust the contrast. The process is simple and intuitive, and takes a a few of minutes at most. An optional mask may be present, which is a bi-level image that indicates which pixels in the image will be subject to the optimization, described next. The mask can significantly speed the optimization process up if lighting changes in some areas of the image are unimportant. Generating the mask can be accomplished through a variety of image processing techniques. Additionally, however, we support automatic generation of the mask through a mode that adds pixels to the mask wherever a user paints on the target image. In pilot studies of the system, many users used this feature extensively.

B. *Optimization*

Given a target image, the optimization finds weights w_i for each down-sampled image that provide the best match to the target image. We take a constrained least-squares approach. Let N be the number of images in the basis set, each of size $m \times n$. We find

$$\min_{w} \|Aw - t\|^2 \tag{1}$$

such that

$$0 \leq w_i \leq 1$$

for all i, where w is the N-dimensional vector of weights, A is an $(m \times n) \times N$ matrix of basis images (that is, each basis image is treated as a vector), and t is the $(m \times n)$ vector representing the target image painted by the user.

The result is a least-squares optimal match to the supplied target image. As the objective function is quadratic, weights for images with weak contributions are rapidly driven to zero. In our experience, the number of significant nonzero weights is consistently small (e.g., 5-15). This greatly reduces the number of images needed for the final lighting solution.

Specific details regarding the optimization process are as follows. We employ a bound-constrained limited-memory variable metric (BLMVM) for constrained optimization [5] that is part of the TAO optimization package [6]. For example, on a target image size of 170 by 227, with 150 basis images, the optimization takes approximately 70 seconds on a 1GHz Pentium. If the user changes the target image, and repeats the optimization, this time using only the images corresponding to the nonzero weights, then the optimization takes less than 1 second, and is interactive. In pilot studies of our system, most users found the initial delay acceptable, and found that they could adequately refine the image as they wanted using the reduced basis set at interactive rates.

The optimization routine we employ is not guaranteed to find a local minimum. However, we tested our optimization function against a simplex-based simulated annealing algorithm with random restarts [23]. For several runs, with a slow annealing schedule so that the optimization took over 30 hours, and with random initial conditions, the simulated annealing algorithm converged close to the same residual cost and with the same set of weights. We also took the solution produced by the BLMVM method and ran simulated annealing with that as an initial condition. The BLMVM method was never improved upon in any test. Therefore, although we have been unable to prove that the above optimization problem has a global solution, our tests indicate that different optimization strategies converge to the same value, giving us confidence in the robustness of our solutions.

After finding the w_i weights, we apply them to the linearized basis images, then re-apply the camera response function to display a preview of the output image. The user then has the option of replacing the target with a grayscale version of this result and can repeat the sketching and optimization cycle until satisfied with the color preview of the output image.

C. Output Assembly

After iterating on the sketching process a few times, users are generally satisfied with the results. In that case, the final image has been created and the process is over. However, in some cases images may contain over-exposed areas where bright specular highlights are present. In some circumstances these specular highlights can be addressed through the use of HDR photographs.

If the gimbal-mounted camera apparatus (*not* the handheld device) was used to capture the image set, then HDR photographs for the final image set can be captured, because the lighting position is known and repeatable. We also assume the object has not moved or can be placed in a repeatable configuration. In this case, we linearize each basis image to remove effects of the camera response curve. As before, we construct a linear output image as a weighted sum of basis images, using the weights determined by the optimization to match the target image. Finally, we re-apply the camera's response function to the linear output image to get the desired result. This HDR technique is not available for the hand-held acquisition method since repeatability is lacking.

D. Issues of Controllability

A natural question to consider is how much control the user has over the lighting, given that the eventual lighting will be a realizable combination of lights with the physical constraints of the mechanisms described in Section IV. To gain a better understanding of this issue, we used Principal Component Analysis (PCA) [13] to compute the eigenimages of the lighting data set [26]. A bright pixel in an eigenimage indicates a pixel of high covariation when considered with other pixels in that eigenimage. Since each eigenimage is a gray-scale image, we can visualize the first three eigenimages by creating a color image with the first eigenimage as the red channel, the second as the green channel, and the third as the blue. Examples of this for various data sets are shown in Figure 6. The resulting similarly colored regions show us areas where, if we change the lighting in that region, the rest of the lighting will covary as well. Thus, it indicates the limits on the fineness of control that we likely have to affect lighting in the image.

The eigenimages give a nice visualization of the three major regions of covariation, but do not give a complete characterization of the controllability, since there are more than three principal components. The question of how many directions of reasonable control there are becomes the classic question of how many principal components to use in dimensionality reduction. The question is not trivial and depends on each data set. Using the *broken stick* model [13] to compute the number of significant directions indicates that there are between four and seven significant lighting covariations for the data sets presented in this paper. The principal values corresponding to the principal components for the data sets shown in Figure 6 are shown in Figure 7. The principal values are normalized by the trace of the covariance matrix; also shown are the corresponding broken stick values that correspond to whether a principal value is significant or not. In all our data sets, the first three principal values account for over 80% of the variability. Of course, if the user is opting for small highlights around edges, then this evaluation may be of limited use.

VI. Results

Every image in Figure 8 shows results from our sketch guided lighting system. Both the moving-head light and the hand-held methods are equally successful at creating arbitrary cleanly-lit images of desktop-sized objects. The data sets gathered by either method are sufficiently dense to allow easy lighting design. Additionally, our system yields reasonable results even when presented with unrealistic targets or highly reflective objects.

Figure 8(a), demonstrates a user interaction sequence with the system. Starting from a uniform grayscale image as the target, the user guides the optimization, iteratively improving the target until she gets the desired output. Figure 8(b) shows how simple approximate sketching on the target image can give an interesting side-lighting effect. Figure 8(c) shows how the highlight can bring out the underlying texture in a surface.

Figure 8(d) shows lighting for a highly specular object. Good lighting for such smooth, highly reflective objects is always difficult, as the light source itself is visible in the reflection. Our system produces results similar to the target image without large objectionable saturated regions. In future systems we may hide the enclosure seams by constructing wide smooth rounded corners resembling a photographers 'cyc.'

Figure 8(f) shows results from the handheld method of Section IV-B. The data gathering time was under 3 minutes, and the results are comparable to the moving-head light method. While the handheld method is not practical for photographing a large collection of objects, it can be an invaluable tool for well-lit photography in the field, or as a simple consumer level implementation.

Figure 9 shows a comparison of a diffuse light source versus a more focused one on both a diffuse and a specular data set. The diffuse object is small model of a castle, which contains interesting areas for self-shadowing. The specular data object is a crystal and gold dragon sculpture. The diffusely light sets were captured for both objects using the hand-held lighting device, and the more focused ones by removing the diffuse reflector and shining the light directly on the models. Target images and masks are shown in column (a). For each target image we used three basis sets; the diffuse (column (b)), the direct (column (c)), and both basis sets together (column (d)). Not surprisingly, the direct result basis set tends to result in more shadow artifacts but sharper highlights than the diffuse basis sets. The combined data set blends facets of both.

Fig. 6. The first three eigenimages of several lighting data sets presented as color images. The first eigenimage is shown in red, the second in blue, and the third in green. Similarly colored regions indicate regions of like covariation, which implies that a lighting change requested by the user in this region will affect the lighting in the rest of the region.

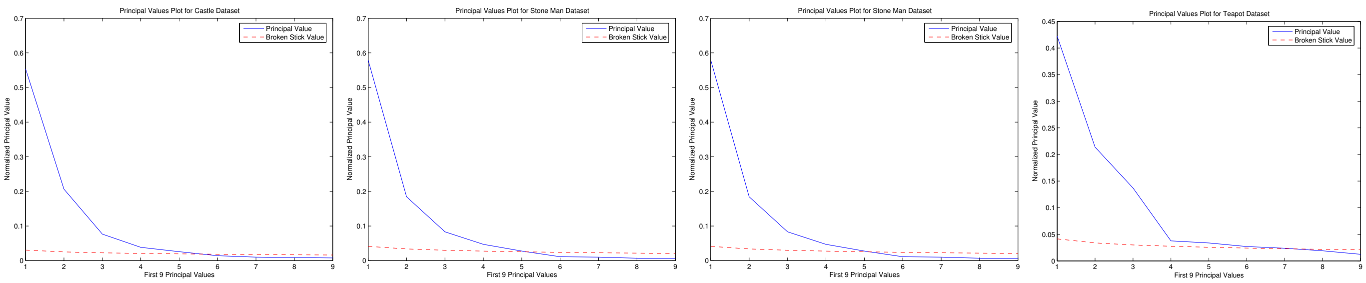

Fig. 7. The normalized principle values (eigenvalues of the covariance matrix divided by its trace) shown together with the broken stick values indicating the number of significant principle components for the data sets shown in Figure 6. The principle value is shown in blue and the broken stick value as a dashed red line.

VII. Discussion and Future Work

The ability to have large area light sources is crucial for photographing highly specular objects. Light source size also affects the sharpness of shadows and highlights. Our system has the unique advantage that larger area light sources can be simulated by combining pictures with overlapping light sources. We could extend our optimization to penalize each distinct light source cluster, thus preventing disjoint highlights. The softness of the light can also be controlled by varying the beam width between a point-source and a large area source as it quickly sweeps over the hemisphere of lighting directions. More advanced moving-head spotlights usually provide controllable spot sizes suitable for this purpose.

Even though the system is aimed primarily at non-professional photographers, a few simple additions can make it a flexible tool for a creative expert to experiment with different lighting designs more easily. For example, the user might specify a simple weighting mask to set the importance of different image regions and influence the optimization process. While weighting masks would make the system more flexible, they would complicate the target sketching process. We do not know yet if the results would warrant the increase in complexity. Also, tools to directly tweak the light position and size on a virtual hemisphere around the object might also aid expert users.

This paper takes the problem of good lighting for desktop photography and finds a simple and practical solution using image-based relighting techniques. More sophisticated image-based measurements might also be achievable while maintaining the simplicity and elegance of the system. For example, we could calibrate our ad-hoc enclosure to measure incident light angles as a function of gimbal angles easily from a set of aiming images of a chrome sphere or ‘light probe’. Combined with surface normals, such calibration might suffice for image-based estimates of BRDF.

Acknowledgments

The authors would like to thank Jingjing Meng and Xing Hu for help in early stages of the project, Kelli Johnson and Nathan Matsuda for help with the construction of the enclosure, Vincent Masselus and Amy Gooch for help with Figure 2, and Elizabeth Seward for help in the pilot studies. This research was funded in part by NSF grants 0238062 and 0237621.

References

[1] A. Agarwala, M. Dontcheva, M. Agrawala, S. Drucker, A. Colburn, B. Curless, D. Salesin, and M. Cohen, “Interactive digital photomontage,” in *SIGGRAPH*, vol. 23, no. 3, 2004, pp. 294–302.

[2] D. Akers, F. Losasso, J. Klingner, M. Agrawala, J. Rick, and P. Hanrahan, “Conveying shape and features with image-based relighting,” in *IEEE Visualization*, 2003.

[3] F. Anrys and P. Dutré, “Image based lighting design,” in *The 4th IASTED International Conference on Visualization, Imaging, and Image Processing*, 2004.

[4] Astron Systems, “Orbiculight lighting system,” http://www.astronsys.com.

[5] S. J. Benson and J. Moré, “A limited memory variable metric method in subspaces and bound constrained optimization problems,” Mathematics and Computer Science Division, Argonne National Laboratory, Tech. Rep. ANL/MCS-P909-0901, 2001.

[6] S. J. Benson, L. C. McInnes, J. Moré, and J. Sarich, “TAO user manual (revision 1.7),” Mathematics and Computer Science Division, Argonne National Laboratory, Tech. Rep. ANL/MCS-TM-242, 2004, http://www.mcs.anl.gov/tao.

(a) Sequence showing successive sketching/optimization iterations to get desired lighting. The first result uses a constant grayscale target, while the others use previous results as starting points for the target image.

(b) Strategic placement of highlights in the target result in an interesting side-lit image.

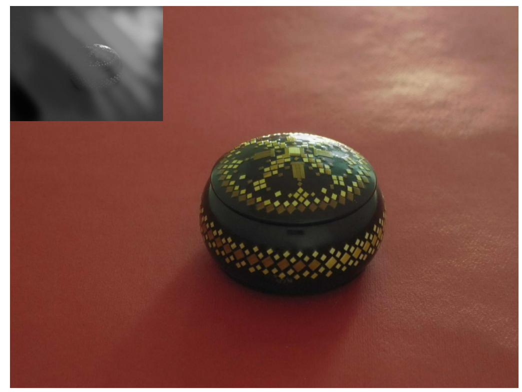

(c) Positioning of highlights reveals underlying texture in the surface.

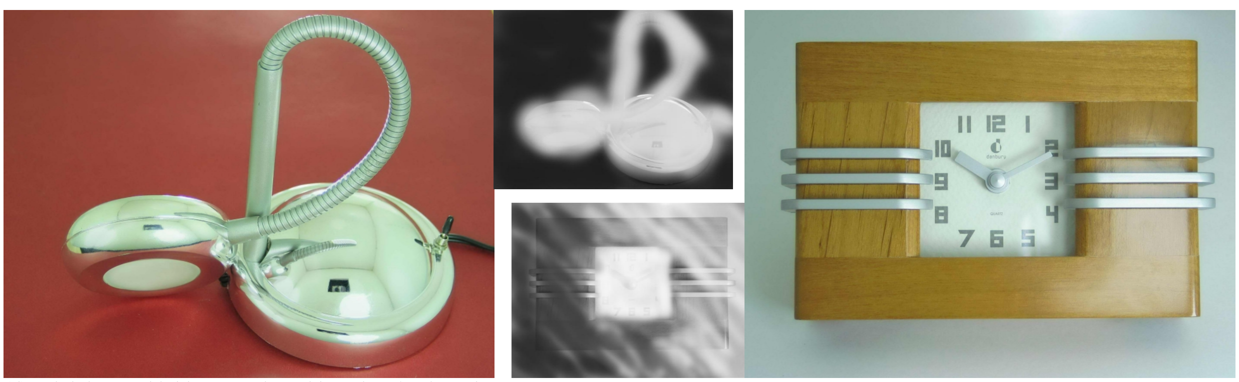

(d) Lighting a highly specular object by forcing the background to be dark.

(e) Target image results in image suggesting illumination from the right.

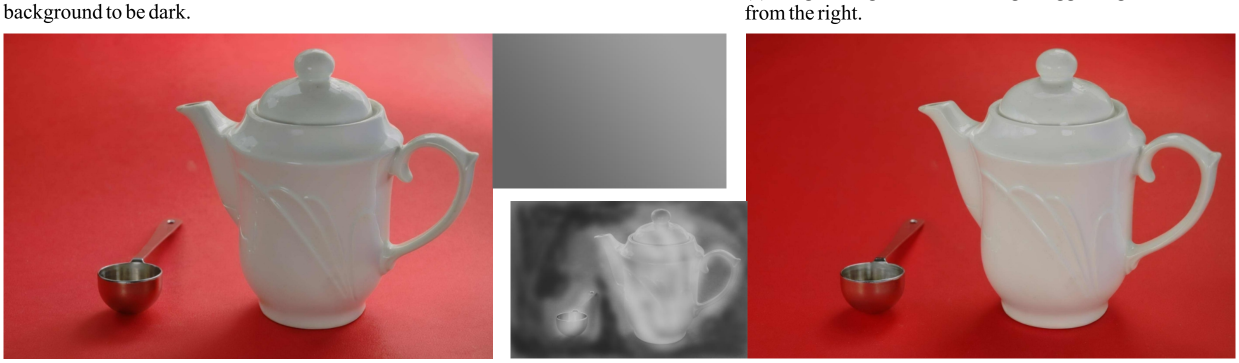

(f) Data captured by the handheld method. Image on the left uses a smooth grayscale gradient as the targe image.

Fig. 8. Sample results from our lighting design system. Each image is is shown with its target image.

Fig. 9. Comparing diffuse versus direct lighting for two different models (view sideways; (a) shows the target image and mask for each set, respectively; (b) shows the diffusely lit set, (c) shows the directly lit set, and (d) shows the image produced from the combined set.

[7] A. C. Costa, A. A. Sousa, and F. N. Ferreira, "Lighting design: a goal based approach using optimisation," in *Proceedings of Eurographics Workshop on Rendering*, vol. 10, 1999, pp. 317–328.

[8] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar, "Acquiring the reflectance field of a human face," in *SIGGRAPH*, 2000, pp. 145–156.

[9] P. Debevec, A. Wenger, C. Tchou, A. Gardner, J. Waese, and T. Hawkins, "A lighting reproduction approach to live-action compositing," in *SIGGRAPH*, 2002, pp. 547–556.

[10] P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *SIGGRAPH*, vol. 31, 1997, pp. 369–378.

[11] M. Fuchs, V. Blanz, and H.-P. Seidel, "Bayesian relighting," in *Rendering Techniques 2005: Eurographics Symposium on Rendering*, K. Bala and P. Dutr, Eds., Konstanz, Germany, July 2005, pp. 157–164.

[12] T. Hawkins, J. Cohen, and P. Debevec, "A photometric approach to digitizing cultural artifacts," in *Proceedings of conference on Virtual reality, archeology, and cultural heritage*, 2001, pp. 333–342.

[13] I. Jolliffe, *Principal Component Analysis*. New York: Springer-Verlag, 1986.

[14] J. K. Kawai, J. S. Painter, and M. F. Cohen, "Radioptimization: goal based rendering," in *SIGGRAPH*, 1993, pp. 147–154.

[15] T. Malzbender, D. Gelb, and H. Wolters, "Polynomial texture maps," in *SIGGRAPH*. ACM Press, 2001, pp. 519–528.

[16] V. Masselus, P. Dutré, and F. Anrys, "The free-form light stage," in *Proceedings of the 13th Eurographics workshop on Rendering*, 2002, pp. 247–256.

[17] V. Masselus, P. Peers, P. Dutré, and Y. D. Willems, "Relighting with 4d incident light fields," in *SIGGRAPH*, vol. 22, no. 3, 2003, pp. 613–620.

[18] MK Digital, "Photo light box," http://www.mkdigitaldirect.com.

[19] J. S. Nimeroff, E. Simoncelli, and J. Dorsey, "Efficient re-rendering of naturally illuminated environments," in *Proceedings of the Fifth Eurographics Workshop on Rendering*, 1994, pp. 359–373.

[20] Photek, "Digital lighthouse," http://www.photekusa.com.

[21] P. Poulin and A. Fournier, "Painting surface characteristics," in *Proceedings of Eurographics Workshop on Rendering*, 1995, pp. 160–169.

[22] P. Poulin, K. Ratib, and M. Jacques, "Sketching shadows and highlights to position lights," in *Proceedings of Conference on Computer Graphics International*, 1997, pp. 56–63.

[23] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, 2nd ed. Cambridge University Press, 1992.

[24] C. Schoeneman, J. Dorsey, B. Smits, J. Arvo, and D. Greenburg, "Painting with light," in *SIGGRAPH*, 1993, pp. 143–146.

[25] R. Shacked and D. Lischinski, "Automatic lighting design using a perceptual quality metric," *Computer Graphics Forum*, vol. 20, no. 3, pp. 215–226, 2001.

[26] M. Turk and A. Pentland, "Face recognition using eigenfaces," *Journal of Cognitive Neuroscience*, vol. 3, no. 1, 1991.